Promotion Price Optimization

This case study shows how traditional price management processes can be improved using machine learning. We describe how machine learning helped to optimize promotions and special offers to increase profits and prevent losses.

Business Problem

Since pricing is directly linked to consumer demand and company profits, pricing decisions are critically important for any business. Even a slightly suboptimal decision-making process inevitably leads to tangible losses, and major mistakes can have grave consequences.

Optimal pricing is a challenging problem for several reasons. One is the complex structure of the price waterfall, which often includes multiple variables such as list prices, discounts, and special offers that need to be optimized. Another reason is the complexity of demand and profit forecasting, which makes it difficult to evaluate new pricing strategies accurately. Finally, nonoptimal pricing decisions are often caused by a lack of coordination between the teams that are responsible for various aspects of pricing.

Grid Dynamics was engaged by a leading apparel retailer that relied heavily on promotions and special offers to drive shopper traffic and incremental sales. The retailer had several distinct levels of promotions, including storewide, category-wide, brand-wide, and product-level promotions. In some cases, particular levels were managed by distinct merchandising teams that lacked coordination.

The main challenge was that the retailer did not have sufficient analytical tools to assess the quality of the final promotion mix and determine how effective that mix was from the standpoint of profitability. A second challenge was to improve the merchandising process by introducing more powerful tools for promotion management and profit forecasting.

Analysis: Environment

The promotion management environment is shown in Figure 9. The merchandiser maintained a database of promotions in which each promotion was configured by triggering rules (e.g., the purchase total must be more than $100) and actions (e.g., provide a $20 discount). For each transaction, the pricing engine pulled active promotions from this database based on products in the shopping cart and then calculated the final sales price applying these promotions.

The distinctive feature of this retailer was a complex set of promotion arbitration rules that were applied by a pricing engine in each shopping cart. The purpose of these rules was to ensure that no promotions that applied to the contents of a shopping cart conflicted with each other and to reconcile overlapping promotions if needed. Since the rules were applied to individual shopping carts, the final sales price of a product was unique for each transaction and could vary depending on other products in the cart. Consequently, the sales price of any given product at any single moment could be described only as a statistical distribution of prices rather than as a single number. This presented a significant challenge to building the price-demand model needed for promotion optimization.

Analysis: Controls

The promotion database contained many promotions, and each promotion had several parameters that could potentially be optimized, including the following:

- Status (enabled or disabled)
- Discount value
- Start and end dates

Figure 9. Promotion price optimization environment.

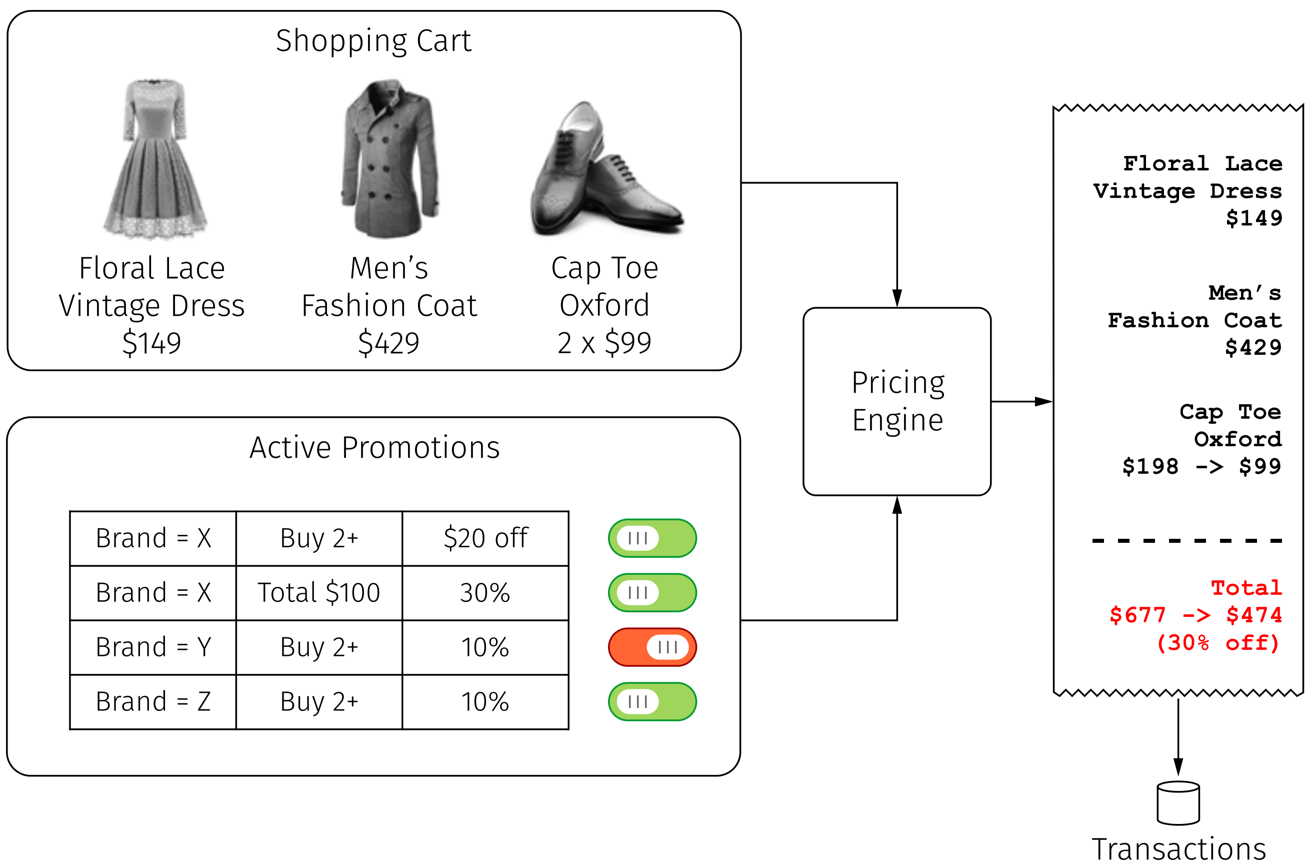

In addition to that, some other parameters of the price waterfall, such as list prices, could potentially be optimized.

Analysis: Objectives

The immediate goal was to optimize the promotion mix by disabling promotions that were likely to cause losses and by enabling ones that increased profits. This was intended to mitigate the problem of nonoptimal promotion management caused by the lack of tools and lack of coordination between teams.

The longer-term goals included the optimization of promotion parameters, such as dates and discount values, as well as automatic promotion-opportunity finding and the creation of new promotions.

The typical optimization objective was to maximize the profit of a given product category by setting optimal promotion parameters. The promotion optimization system, however, also had to forecast future demand (number of sold units) and revenue so that these numbers could be used in planning.

Finally, the optimization could be the subject of several constraints. For example, the ability to accelerate demand using promotions could be limited by inventory constraints and by the risk of harming profits by running out of stock. Another example is clearance sales, in which merchandise has to be sold by a certain date and demand-accelerating parameters have to be set accordingly.

Solution: Vision

We envisioned a solution involving a fully automated system that could optimize the promotion mix based on the objective of profit maximization. The first step was to create a tool that imported promotions

Figure 10. Solution vision.

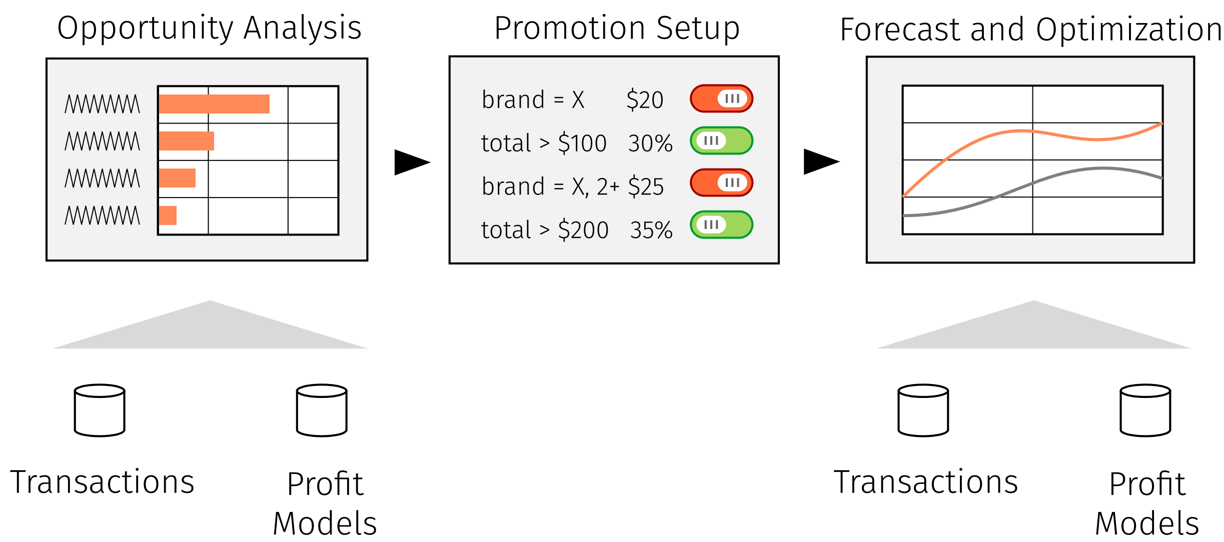

from the legacy promotion-configuration system and allowed a merchandiser to do what-if analysis of the promotion mix. This functionality corresponds to the right side of the workflow in Figure 10. The merchandiser had to be able to turn individual promotions on and off using simple toggles, and the system had to forecast the expected revenue, profit, and quantity sold based on this configuration and historical data in near real time.

The next phases of development included additional features shown in Figure 10. One of these features was automatic mix optimization, so that the merchandiser did not have to search manually for the optimal promotion combination. The second important feature was the opportunity finder, which enabled the system not only to optimize the existing mix but also to create new promotions automatically.

Solution: Architecture and AI Usage

Price optimization generally relies on the ability to predict demand or profit as a function of price-related parameters, such as list prices, discounts, and promotions. The demand- or profit-prediction model, in turn, can be created by learning patterns and dependencies from the historical data. Once the model is created, the optimization system can use it to evaluate diverse possible values of the parameters and determine the optimal combination. Consequently, modeling profit and demand was the key component of the solution.

The profit modeling subsystem is shown on the left side of Figure 11. It includes data collection components that consolidate transactional data, product catalog data, promotion calendars, and other signals that can potentially help to predict demand and profit. The data science team synthesized various derivative features from this data and then trained models that predict the demand or profit time series. The promotion optimization server uses these models to evaluate promotion combinations configured by the merchandiser in the user portal.

As we already mentioned, one of the challenges was that, due to dynamic promotion-arbitration rules, the products did not have fixed sales prices. Consequently, it was not possible simply to set promotional markdowns and evaluate the profit prediction model for them, even if such a model was available. To work around this problem, we developed a pricing simulator that used shopping cart statistics to estimate the distribution of sales prices for any given product. This simulator analyzed a large number of orders to determine products that were frequently bought together and then used

Figure 11. Solution design. Components that use AI capabilities are marked with a star.

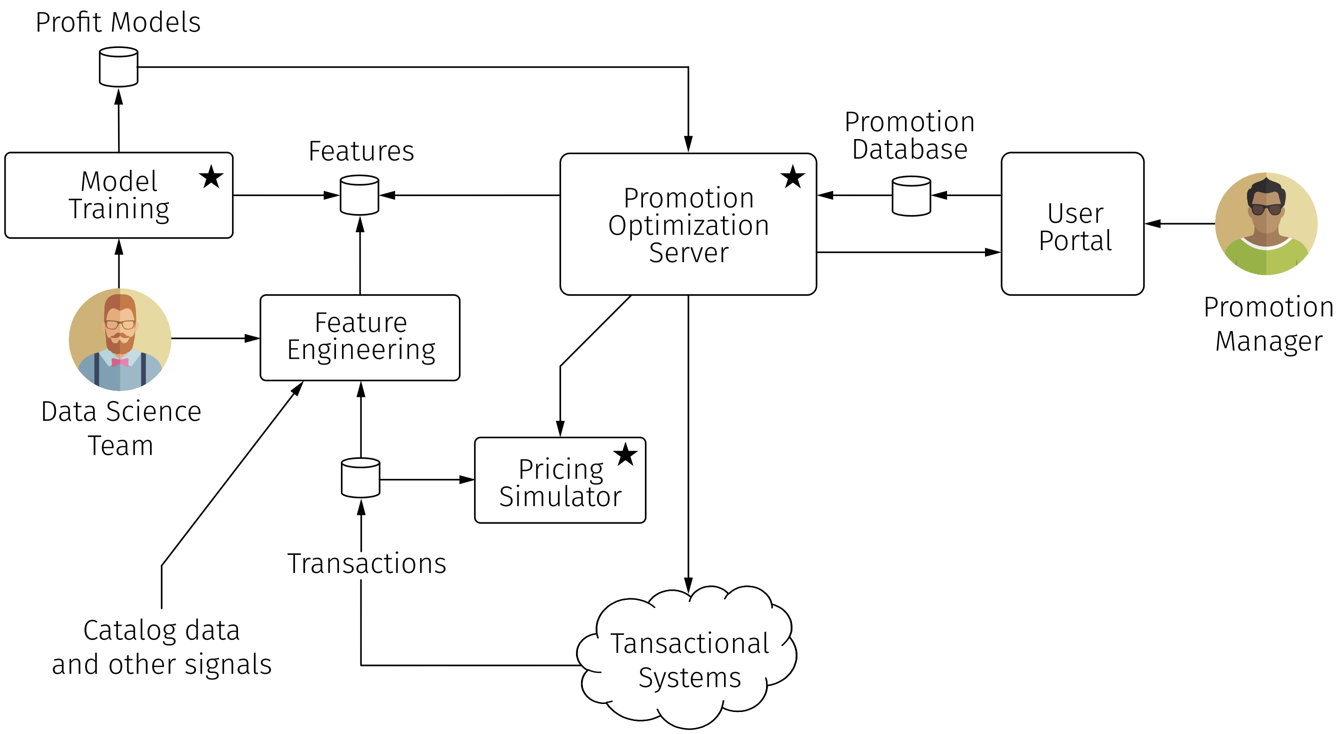

that data to estimate possible product sales prices given a new promotion setup. The optimization engine uses the simulator to estimate this price distribution for various promotion setups.

Results

- One of the main results of the program was a transformation of the traditional promotion management process into a new process that uses scientific methods for promotion optimization. The new solution brings the power of data science and economic modeling to daily business operations.
- The new method of promotion management provides certain guarantees that pricing decisions are near optimal; the old process provided no such guarantees.
- The solution helped to detect more than a thousand products that systematically caused losses when they were promoted. These products accounted for about 90% of all losses caused by nonoptimal promotion setup.
- The solution was able to accurately forecast profits, revenues, and demand three months ahead.

Related Business Cases

The described solution can be applied to a number of business cases that require demand, profit, or revenue forecasting. This includes a large number of price-optimization use cases, including optimization of markups, markdowns, clearance-sale events, and flash sales. It is especially useful for price segmentation and dynamic and personalized pricing when pricing decisions need to be done at a high level of granularity (e.g., store level or customer-segment level) or frequently need to be re-optimized.

Another group of business cases is related to inventory management. Use cases such as stock level optimization also rely on accurate demand prediction, and the described solution can be integrated with inventory management systems to optimize certain decisions.